



**Eur päisches
Patentamt**

**European
Patent Office**

**Office eur péen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03368003.4

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

DEN HAAG, DEN
THE HAGUE, 27/02/03
LA HAYE, LE



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.:
Demande n°: 03368003.4

Anmeldetag:
Date of filing:
Date de dépôt: 07/01/03

Anmelder:
Applicant(s):
Demandeur(s):
INTERNATIONAL BUSINESS MACHINES CORPORATION
Armonk, NY 10504
UNITED STATES OF AMERICA

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:

A method and system for monitoring performance of distributed applications

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:

/

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing:
Etats contractants désignés lors du dépôt:

AT/BG/BE/CH/CY/CZ/DE/DK/EE/ES/FI/FR/GB/GR/HU/IE/IT/LI/LU/MC/

Bemerkungen:
Remarks:
Remarques:

A METHOD AND SYSTEM FOR MONITORING PERFORMANCE OF DISTRIBUTED APPLICATIONS

Technical field

The present invention relates to the data processing
5 field, and more specifically to a method and a corresponding
system for monitoring performance of distributed applications.

Background art

Distributed applications have become increasingly popular
10 in the last years, particularly following the widespread
diffusion of the Internet. In a distributed application,
client computers access resources managed by server computers
through a network. A typical example is that of an e-business
application, wherein a user may download a login page, fill-in
15 a form with his/her username and password, and then receive
information (for example, about a personal bank account) from
the server.

Tools for monitoring performance of distributed
applications play a key role in their management.
20 Particularly, a system administrator can get instantaneous
notification when a user is experiencing any problem (so that
appropriate steps can be taken to remedy the situation);
alternatively, the collected information can be logged and
accurate counts tracked over time. For example, the
25 information provided by the monitoring performance tools is
essential for service level agreements or for threshold and/or
availability monitoring; moreover, the same information is
very useful to measure workloads for capacity planning and
charge-back accounting.

However, these tools (like any measurement systems) inevitably interfere with the quantities under measure; therefore, the correct tuning of the performance monitoring tools is of the utmost importance, in order to avoid adversely
5 affecting operation of the whole system.

A solution known in the art for monitoring performance of distributed applications is provided by the Application Response Measurement (ARM) standard, as described in "The Application Response Measurement (ARM) API, Version 2", Mark
10 W. Johnson, Tivoli Systems, December 1997 (available at <http://regions.cmg.org/regions/cmgarmw/index.html>). The standard defines some API calls, which can be used to ask an agent to measure transactions and to make the information available to management applications. In this way, an accurate
15 picture of the actual workload of the system can be obtained.

The ARM standard also supports the use of correlators, which provide child/parent information needed to trace how transactions and corresponding sub-transactions relate to each other. The correlators are very useful to breakdown the
20 complexity of the distributed application, so as to facilitate the analysis of the collected information. For example, when a transaction is slow it is possible to know which sub-transaction(s) contribute most to the delays.

A distributed application must be correctly instrumented
25 for monitoring its performance using the ARM standard. First of all, this procedure requires the identification of the key transactions to be monitored. The distributed application is then modified by embedding calls to the ARM APIs where necessary.

30 However, the solution described above is very rigid since the key transactions must be defined statically. The cited document only suggests a technique for exploiting the format of the correlators so as to use the tracing selectively (for example, when the response time of a client begins to be
35 unacceptable). However, once the calls to the ARM APIs have been inserted into the distributed application, no way is

provided for controlling the transactions to be monitored dynamically; conversely, any change requires the updating of the corresponding source code and its deployment to the different (client and server) computers where the distributed
5 application is running.

Therefore, a wrong selection of the key transactions can be detrimental to the operation of the whole system. Particularly, when few transactions are selected the collected information may be useless; conversely, monitoring a great
10 number of transactions may result in application delays and system overhead.

Moreover, the instrumentation of the distributed application is not a tenable option when its source code is not available. This drawback is particular acute for
15 pre-loaded or packaged applications; a typical example is that of the browsers that are installed on millions of clients for accessing the Internet.

A different solution for monitoring performance of distributed applications where source code changes are not
20 possible is described in "Service management using the application response measurement API without application source code modification", Martin Haworth, Resource and Performance Management Solutions Network and System Management Division, Hewlett-Packard Company, June 1997 (on
25 <http://regions.cmg.org/regions/cmgarmw/index.html>). This article proposes capturing a script by recording the user actions on the client, for example, by means of a Remote Terminal Emulation (RTE) technique. The user script is edited to include calls to the ARM APIs for each desired transaction.
30 The user script can then be scheduled to run at an appropriate interval.

However, the proposed technique only provides an emulation scenario, wherein the performance parameters that are measured are always artificial in nature. Therefore, this solution is
35 very limited in that it cannot provide an accurate picture of the performance of the real transactions.

Summary of the invention

It is an object of the present invention to provide a method and a corresponding system for monitoring performance of distributed applications, which method and system support a dynamic control scheme.

It is another object of the present invention to provide a simple and flexible solution for selecting the transactions to be monitored.

10 It is yet another object of the present invention to facilitate the tuning of the performance monitor process.

The accomplishment of these and other related objects is achieved by a method of monitoring performance of distributed applications including the steps of: a client computer
15 originating a request of service for a server computer, if the request meets at least one predefined condition, enabling the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and updating the request by inserting a correlation identifier,
20 transmitting the request to the server computer, if the request includes the correlation identifier, enabling the measuring on the server computer of the at least one performance parameter for a sub-transaction originating from the request, executing the sub-transaction, associating the
25 correlation identifier with the at least one performance parameter measured on the server computer, and associating the correlation identifier with the at least one performance parameter measured on the client computer.

The present invention also provides different computer
30 programs for performing the method, together with corresponding products storing the programs. Furthermore, a data processing system for monitoring performance of distributed applications, a client computer and a server computer for use in the system are also encompassed.

35 The novel features believed to be characteristic of this invention are set forth in the appended claims. The invention itself, however, as well as these and other related objects

and advantages thereof, will be best understood by reference to the following detailed description to be read in conjunction with the accompanying drawings.

Brief description of the drawings

- Figure 1a** is a schematic block diagram of a data processing system in which the method of the invention can be used;
- Figure 1b** depicts the functional blocks of a generic computer of the system;
- Figure 2** shows a partial content of the working memories of a client and of a server in the system;
- Figures 3a-3d** illustrate an activity diagram describing the execution of a transaction.

5

Detailed description of the preferred embodiment

With reference in particular to Figure 1a, a data processing system 100 with a distributed architecture based on the Internet is shown. The Internet is formed by millions of
10 computers, which are connected to each other through a network infrastructure 105; each computer is uniquely identified by a corresponding IP address. Client computers 110 get into the Internet through associated Internet Access Providers (ISPs) 115; access to the Internet allows users of the clients 110 to
15 exploit shared resources supported by server computers; for example, the users can exchange information, send and receive e-mails, and view documents. A particular subsystem of servers 120 (the World Wide Web) manages hypertext documents, known as web pages; each web page is formatted in the HTML, a language
20 that supports links to other documents, as well as graphics, audio, and video files. The web uses the HTTP protocol, which

defines how messages are structured and transmitted, and what actions the clients 110 and the servers 120 should take in response to various commands.

As shown in Figure 1b, a generic computer of the system 5 (client or server) is formed by several units that are connected in parallel to a communication bus 130. In detail, a microprocessor (μ P) 135 controls operation of the computer, a Random Access Memory (RAM) 140 is directly used as a working memory by the microprocessor 135, and a Read Only Memory (ROM) 150 stores basic code for a bootstrap of the computer. Several peripheral units are further connected to the bus 130 (by means of respective interfaces). Particularly, a mass memory consists of a magnetic hard-disk 155 and a driver 160 for reading CD-ROMs 165. Moreover, the computer includes input devices 170 (for 15 example, a keyboard and a mouse), and output devices 175 (for example, a monitor and a printer). A network Interface Card (NIC) 180 is used to connect the computer to the network infrastructure.

However, the concepts of the present invention are also 20 applicable when the system has another architecture (for example, based on a Local Area Network or LAN), or when each computer has a different structure or includes other units. Similar considerations apply if the shared resources are managed by a single server, if the system includes nodes 25 operating alternatively either as clients or as servers, if proxy computers and/or gateway computers are provided, and the like.

Moving now to Figure 2, a partial content of the working memories of a generic server and of a generic client is shown; 30 the information (programs and data) is typically stored on the respective hard-disks and loaded (at least partially) into the working memories when the programs are running, together with an operating system and other application programs (not shown in the figure). The programs are initially installed onto the 35 hard disks from CD-ROMs.

A browser 205 allows a user of the client to surf the Internet, in order to download and display web pages from selected servers. Each web page is identified by a global

name, known as Uniform Resource Locator (URL). The URL consists of a first part indicating the protocol to be used (for example, the HTTP) and a second part actually indicating the requested web page; the web page is selected by means of a domain name (identifying one or more IP addresses), possibly followed by the path to the requested web page.

Any request of the user causes the execution of a transaction. The transaction consists of a sequence of information exchanges and related operations that are required to satisfy the request; the transaction ends with the return of a response to the user. The transaction must be completed substantially immediately, and in any case in a reasonable amount of time to allow an interaction of the user with the system.

The browser 205 generates a HTTP request for the transaction, which is redirected to a sensor 210. The sensor 210 controls a monitoring agent 215. The monitoring agent 215 measures one or more performance parameters of the transaction (for example, its duration) directly on the client; moreover, the monitoring agent 215 generates a correlator for the transaction, which is returned to the sensor 210. The measured parameter of the transaction is stored in a log 220 (together with its correlator). The sensor 210 further accesses a filter 225, which stores a predefined pattern of domain names to be monitored.

The HTTP request is modified by the sensor 220 (if necessary) inserting the correlator. The modified HTTP request is then returned to the browser 205, in order to be transmitted to the corresponding server. For this purpose, the browser 205 interfaces with a networking module 230, which processes a set of protocol layers working together for defining communication over the Internet. The networking module 230 is also coupled with the monitoring agent 215, in order to transmit the log 220 periodically (for example, every night) to a dedicated server acting as a collector of the measured parameters.

The networking module 230 establishes a connection with a corresponding module 235 running on the server to which the

(modified) HTTP request is directed. The HTTP request is processed by a web server module 240. For this purpose, the web server 240 accesses a repository 245 of web pages; the web pages may be either static (when their content cannot be
5 changed) or dynamic (when their content is defined at run-time).

Each HTTP request received by the web server 240 causes the execution of a corresponding sub-transaction (involving the requested web page to be fetched and returned to the
10 client through the corresponding connection). Moreover, the sub-transaction may require one or more additional sub-transactions, which are performed under the control of corresponding applications 250. For example, the applications 250 retrieve information requested by the user from a database
15 stored on the server. In addition, some sub-transactions may require the transmission of one or more additional HTTP requests to auxiliary servers (through the networking module 235); each (auxiliary) HTTP request in turn causes the
20 execution of one or more sub-transactions, which end with the return of a response to the (main) server. The same procedure may be repeated recursively on the auxiliary servers once or more times.

The web server 240 and the applications 250 control a monitoring agent 255, which measures one or more performance
25 parameters of each sub-transaction directly on the server. The measured parameters are stored in a log 260, together with the correlator of the parent transaction (received from the client in the HTTP request). In addition, if the sub-transaction involves the execution of nested sub-transactions, the
30 monitoring agent 255 generates a further correlator; this correlator is associated with the measured parameters for the sub-transaction in the log 260. Moreover, the correlator is returned to the application 250, which updates the corresponding auxiliary HTTP request accordingly. The
35 monitoring agent 255 is also coupled with the networking module 235, in order to transmit the log 260 periodically to the collector server.

Preferably, the monitoring agent 215 (on the client) and the monitoring agent 255 (on the server) support API calls conforming to the ARM standard. The ARM standard (Version 2.0) defines six APIs. Two APIs `arm_init` and `arm_getid` are typically executed only once when the module (sensor or web server) exploiting the monitoring agent initializes. The API `arm_init` is called passing the name of the module and returns a unique identifier thereof; the API `arm_getid` is called passing the identifier of the module and the name of a transaction class, and returns a unique identifier of the transaction class. A further API `arm_end` (receiving the identifier of the module) may be used to signal to the monitoring agent that the module is shutting down. Two APIs `arm_start` and `arm_stop` are used to monitor an actual transaction. Particularly, the API `arm_start` indicates that the transaction has begun execution; the API `arm_start` is called passing the identifier of the transaction class and returns a unique handle for the transaction. The API `arm_stop` instead indicates that the transaction has been completed; the API `arm_stop` is called passing the handle of the transaction and a value denoting its status (good, error or abort). A further API `arm_update` is available for signaling that a very long transaction (identified by the handle passed as a parameter) is still active.

The API `arm_start` may be called requesting the monitoring agent to return a correlator for the transaction. The correlator includes a flag, which is used to enable the tracing of the transaction; moreover, the correlator includes the handle of the transaction and an identifier of the computer on which the monitoring agent is running (for example, its IP address). On the same API `arm_start`, the module may also provide the correlator of a parent transaction. In this way, the correct child/parent relationship among different transactions can be easily traced.

For example, let us assume that a client A starts a transaction T1 requesting a correlator, which is assigned C1. The client A sends a corresponding request to a server B, and

includes C1 in the request. In response thereto, the server B starts a sub-transaction T2, passing C1 as the parent correlator; at the same time, a further correlator for the sub-transaction T2 is requested and assigned C2. In turn, the
5 server B sends a further request to a server C, and includes the correlator C2 in the request. The server C starts an auxiliary sub-transaction T3, passing C2 as the parent correlator. The total picture of the transaction can be reconstructed knowing that the transaction T1 is the parent of
10 the sub-transaction T2 (via the correlator C1), and that the sub-transaction T2 is the parent of the sub-transaction T3 (via the correlator C2).

However, the concepts of the present invention are also applicable when the transaction originates different requests
15 of service (each one involving the execution of alternative operations), when the HTTP request is intercepted with another technique (for example, using a hooking technology), or when other performance parameters are measured; alternatively, the ARM calls include additional information about the transaction
20 (for example, its internal congestion, the level of resources available, or diagnostic data), or the correlators are replaced with equivalent identifiers; moreover, the monitoring agents may support calls conforming to a different version of the ARM standard or even calls conforming to another
25 specification. Similar considerations apply if the programs are provided on any other computer readable medium (such as one or more floppy-disks), if the programs and data are structured in a different manner, if alternative modules or functions are provided, and the like. For example, the parameters measured
30 on every computer are aggregated before being logged, or the logs are collected in a different way (or they are simply available for viewing without being transmitted periodically to the collector server); alternatively, the monitoring agents generate an alarm when significant problems occur, or a module
35 is provided for identifying and solving the problems automatically.

Considering now Figures 3a-3d, a process 300 corresponding to the execution of a transaction begins at the

black start circle 302 in the swim-lane of the browser. Proceeding to block 303, a web page requested by the user is downloaded on the client; the web page is then displayed on the monitor of the client at block 304. As soon as the user selects
5 a link on the current web page (for example, clicking with the mouse on its graphical representation), a corresponding HTTP request for the corresponding server is generated at block 305.

The HTTP request is intercepted by the sensor at block 306. The process then branches at block 307, wherein a test is made
10 to determine whether the HTML code defining the current web page includes a predetermined keyword for enabling the monitoring of the transactions originating therefrom. If the keyword has been found, a correlation flag is asserted at block 308; conversely, the correlation flag is deasserted at block 310. In both cases,
15 the process merges again at block 312.

Continuing to block 314, the sensor verifies whether the domain name of the server (specified in the HTTP request) matches the filtering pattern. If not, the correlation flag is deasserted at block 316, and the process continues to block 318;
20 conversely, the flow of activities descends into block 318 directly.

Considering now decision block 318, if the correlation flag is deasserted the process continues to block 320 (described in the following). Conversely, a test is made at block 322 to
25 determine whether an additional filter at the level of the link is enabled. If so, the sensor verifies at block 324 whether the method specified in the HTTP request (corresponding to the definition of the link in the HTML code) includes a predetermined keyword for enabling the monitoring of the
30 transactions originating from the link. If the keyword has not been found, the process continues to block 320. If the filter at the level of the link is not enabled (block 322) or the HTTP request includes the keyword (block 324) the process passes to block 326. In this case, the sensor requests the monitoring
35 agent to start measuring the duration of the transaction resulting from the HTTP request (invoking the API arm_start). The API arm_start returns a handle of the transaction and a corresponding correlator at block 328. Continuing to block 330,

the HTTP request is updated inserting the correlator. The flow of activity then passes to block 320.

With reference now to block 320, the HTTP request (possibly updated) is sent to the corresponding server; in this way, the
5 (parent) correlator may be provided to the server without requiring any additional HTTP message. In response thereto, the web server verifies at block 332 whether the HTTP request includes the parent correlator. If so, the process enters block 334, wherein the web server requests the monitoring agent
10 (running on the server) to start measuring the duration of the sub-transaction resulting from the HTTP request. For this purpose, the API arm_start is called passing the parent correlator received from the client; the API arm_start returns a handle of the sub-transaction and a corresponding correlator.

15 Execution of the sub-transaction is started at block 336. Let us assume that the sub-transaction at block 338 involves the call of a query on a selected application running on the server. If the parent correlator had been included in the HTTP request (decision block 340), the call is updated at block 342 adding
20 the correlator of the sub-transaction as a further parameter.

In any case, the process then continues to block 344, wherein the query is called on the application (possibly passing the correlator). In response thereto, the process proceeds to block 346 in the swim-lane of the application; assuming that the
25 query involves the execution of a further sub-transaction on an auxiliary server, the application generates a corresponding HTTP request. A test is then made at decision block 348 to verify whether the correlator has been passed to the application in the call. If so, the process enters block 350 wherein the auxiliary
30 HTTP request is updated adding the correlator; the process then continues to block 352. Conversely, the flow of activities descends into block 352 directly.

Moving to block 352, the auxiliary HTTP request (possibly updated) is sent to the corresponding server. In response
35 thereto, a test is made at decision block 354 (in the swim-lane of the auxiliary server) to determine whether the auxiliary HTTP request includes the (parent) correlator. If so, the process enters block 358 wherein the monitoring agent (on the auxiliary

server) is requested to start measuring the duration of a further sub-transaction originating from the auxiliary HTTP request; the process then continues to block 360. Conversely, the flow of activities descends into block 360 directly.

5 The operations involved by the auxiliary HTTP request are then executed at block 360. As soon as the corresponding sub-transaction has been completed, if the parent identifier had been included in the auxiliary HTTP request (decision block 362) the blocks 364-366 are executed; in any case, the process then
10 continues to block 368 (described in the following). Particularly, at block 364 the monitoring agent is requested to stop measuring the duration of the auxiliary sub-transaction; for this purpose, the API arm_stop is called passing the handle of the sub-transaction (returned by the API arm_start). The
15 measured parameter is then logged at block 366, together with the parent correlator.

Considering now block 368, the result of the sub-transaction executed on the auxiliary server is returned to the application running on the (main) server. The process
20 continues to block 370, wherein the application in turn returns this result to the web server. With reference now to block 372 in the swim-lane of the web server, a dynamic web page including the returned information is generated. If the parent correlator had been included in the HTTP request received from the client
25 (decision block 374), the blocks 376-378 are executed; in any case, the process then continues to block 380 (described in the following). Particularly, at block 376 the monitoring agent is requested to stop measuring the duration of the sub-transaction (calling the API arm_stop and passing the handle of the
30 sub-transaction); the measured parameter is then logged at block 378, together with the correlator and its parent correlator.

With reference now to block 380, the web page so generated is returned to the browser. In response thereto, the web page is displayed on the monitor of the client at block 382 (in the
35 swim-lane of the browser). As soon as the transaction has been completed, the process continues to decision block 384 in the swim-lane of the sensor. If the correlator had been included in the HTTP request, the blocks 386-388 are executed. Particularly,

at block 386 the monitoring agent is requested to stop measuring the duration of the transaction (calling the API `arm_stop` and passing the handle of the transaction); the measured parameter is then logged at block 388, together with its correlator. In
5 any case, the process ends at the concentric white/black stop circles 390.

For example, the current web page displayed on the client is defined by the following HTML code:

```
<HTML>
10 <HEAD>
    <META name="Correlation" content="Enabled">
    ...
</HEAD>
<BODY>
15 <A HREF="http://MyDomainName/MyDestinationPage?Correlation=Enabled">
    Click here</A>
</BODY>
</HTML>
```

20 The HTML code starts with the `<HTML>` tag and ends with the `</HTML>` tag. The definition of what the web page is about is put in a header between the `<HEAD>` and `</HEAD>` tags. All the information to be included in the web page fits in a body between the `<BODY>` and `</BODY>` tags. In the example at issue,
25 the header includes a meta field called "Correlation"; the meta field, when enabled, specifies that the transactions originating from the web page must be monitored. Moreover, the web page includes a link to an anchor destination defined by the following URL:

30 "http://MyDomainName/MyDestinationPage"
(wherein "MyDomainName" is the domain name and "MyDestiantionPage" is the location of the requested web page). The parameter "?Correlation=Enabled" added to the URL specifies that the specific transaction originating from the
35 link must be monitored.

When the user selects this link, the following HTTP request is generated:

```
GET <SP>
http://MyDomainName/MyDestinationPage?Correlation=Enabled <SP>
HTTP_Version <CRLF>
<CRLF>
5 <CRLF>
```

The HTTP request consists of header fields and an entity body, which are separated by a null line (CRLF). The header fields specify that the method GET is requested to retrieve the web page identified by the URL:

```
10 "MyDomainName/MyDestinationPage"
(the parameter ?Correlation=Eanbled" is discharged by the web
server).
```

If the domain name "MyDomainName" matches the filtering pattern, the HTTP request is modified inserting a new field

```
15 "CORR" specifying the value of the correlator (for example,
"MyCorrelator"):
```

```
GET <SP>
http://MyDomainName/MyDestinationPage?Correlation=Enabled <SP>
HTTP_Version <CRLF>
20 CORR: "MyCorrelator" <CRLF>
<CRLF>
<CRLF>
```

Likewise, the following HTML code defines a current web page that is used to submit a form to the web server:

```
25 <HTML>
<HEAD>
<META name="Correlation" content="Enabled">
...
</HEAD>
30 <BODY>
<FORM Method="POST"
Action="http://MyDomainName/MyDestinationPage">
<INPUT Type="HIDDEN" Name="Correlation" Value="Enabled">
```

```
<INPUT Type="SUBMIT" Value="Click Here">
</FORM>
</BODY>
</HTML>
```

5

In this case, a hidden tag (Name="Correlation" and Value="Enabled") is used to specify that the transaction originating from the link must be monitored.

When the user selects this link, the following HTTP request
10 is generated:

```
POST <SP>
http://MyDomainName/MyDestinationPage <SP>
HTTP_Version <CRLF>
<CRLF>
```

15 <CRLF>

The HTTP request is then modified (assuming that the domain name "MyDomainName" matches the filtering pattern) by inserting the field "CORR" (with the corresponding value "MyCorrelator"):

```
POST <SP>
20 http://MyDomainName/MyDestinationPage <SP>
HTTP_Version <CRLF>
CORR: "MyCorrelator" <CRLF>
<CRLF>
<CRLF>
```

25 The performance parameters measured with the process described above are then available to a system administrator (for example, after being consolidated on the collector server). Particularly, the correlators associated with every measured parameter are used to trace the correct child/parent
30 relationship among the different transactions. For example, when a problem is experienced by the user a first level of detail allows the system administrator to identify the component of the system (client, server or network) causing the problem. The information relating to the corresponding sub-transactions can

then be expanded in a top-down manner, until the true cause of the problem is precisely identified.

However, the concepts of the present invention are also valid when an application (consisting of the programs on each client and the programs on each server) performs an equivalent method. Similar considerations apply if the correlator is transmitted to the server in a different manner, if the web pages are replaced with equivalent documents, or if different (general and/or local) identifiers are used in the definition of the web page to enable the monitoring of the transactions. Alternatively, the monitoring agents are controlled in another way, the duration of one or more sub-transactions (under the control of corresponding applications) is measured on every server, the transaction involves multiple requests to different servers, and the like.

More generally, the present invention proposes a method of monitoring performance of distributed applications. The method starts with a client computer originating a request of service for a server computer. If the request meets at least one predefined condition, the measuring on the client computer of one or more performance parameters for a transaction corresponding to the request is enabled, and the request is updated by inserting a correlation identifier. The method continues transmitting the request to the server computer. If the request includes the correlation identifier, the measuring on the server computer of the performance parameter for a sub-transaction originating from the request is enabled. The sub-transaction is then executed. The method now involves associating the correlation identifier with the performance parameter measured on the server computer; moreover, the correlation identifier is also associated with the performance parameter measured on the client computer.

The method of the invention makes it possible to control the monitoring of distributed applications dynamically. It should be emphasized that this result is achieved directly filtering the actual requests that are originated on every client computer.

The devised solution provides a simple and flexible way for selecting the transactions to be monitored.

As a consequence, the tuning of the performance monitoring process is strongly facilitated. In this way, the transactions can be monitored selectively only when necessary; it is then possible to collect valuable information without substantially increasing the run-time overhead of the system.

The preferred embodiment of the invention described above offers further advantages.

10 Particularly, every request is intercepted by a module running on the client (which verifies whether the request meets the predefined conditions).

In this way, the process is completely opaque to the module originating the requests on the client computer; this feature is very useful when the source code of this module (for example, the browser installed on the client) cannot be changed.

Particularly, the same operations are recursively executed on one or more auxiliary servers.

20 The proposed solution is particularly advantageous when the transaction involves several interacting tiers working on different computers.

However, the solution according to the present invention leads itself to be implemented even updating the source code of the distributed application, or with transactions that are executed on a single server.

In a particular embodiment of the invention, the measuring is enabled when the address of the server matches a predefined pattern stored on the client.

30 This method can be used when all the requests directed to a specific set of servers have to be monitored.

In addition or in alternative, the measuring is enabled according to the document currently displayed on the client.

The proposed filtering scheme can be controlled centrally from the server, and it is immediately effective on every client requesting the document (without any intervention on the same).

A suggested choice for implementing this filtering scheme is to enable the measuring when a definition of the document includes a global enabling identifier.

This method can be used when the server provides a number
5 of services, and only some of them need to be monitored.

In a different embodiment of the invention, the measuring is enabled when a definition of the link includes a local enabling identifier.

The proposed method provides a very selective filtering
10 scheme.

However, the method according to the present invention leads itself to be implemented only with one or more of the above-described filtering schemes, or even with filters based on different criteria (for example, a geographical location of
15 the client). Alternatively, the addresses, documents and/or links are filtered with different algorithms; for example, the servers to be monitored are selected by the access provider (even according to their IP addresses), or an additional filtering pattern is stored on the client for the URL of the
20 documents to be monitored.

Advantageously, the solution according to the present invention is implemented with a computer program application, which is provided as a corresponding product stored on a suitable medium; the application consists of one or more
25 programs installed on each client and one or more programs installed on each server. However, it should be noted that the different programs running on the clients or on the servers, and even the programs used on the clients to intercept the requests of service are suitable to be implemented and put on the market
30 as stand-alone products (in order to be integrated into pre-existing systems).

Alternatively, each program is pre-loaded onto the hard-disk, is sent to the respective computer through the Internet, is broadcast, or more generally is provided in any
35 other form directly loadable into a working memory of the computer. However, the method according to the present invention leads itself to be carried out with an application having a

different architecture, or even with a hardware structure (for example, integrated in a chip of semiconductor material).

Naturally, in order to satisfy local and specific requirements, a person skilled in the art may apply to the
5 solution described above many modifications and alterations all of which, however, are included within the scope of protection of the invention as defined by the following claims

CLAIMS

1. A method (300) of monitoring performance of distributed applications including the steps of:

5 a client computer originating (303-305) a request of service for a server computer,

if the request meets (307-324) at least one predefined condition, enabling (326) the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and updating (328-330) the

10 request by inserting a correlation identifier,

transmitting (320) the request to the server computer,

if the request includes (332) the correlation identifier, enabling (334) the measuring on the server computer of the at least one performance parameter for a sub-transaction

15 originating from the request,

executing (336-372) the sub-transaction,

associating (378) the correlation identifier with the at least one performance parameter measured on the server computer, and

20 associating (388) the correlation identifier with the at least one performance parameter measured on the client computer.

2. The method (300) according to claim 1, wherein the step of originating (303-305) the request is performed under the
25 control of a first module (205) running on the client computer, the method further including the step of intercepting (306) the request by a second module (210) running on the client computer, the step of enabling (326) the measuring of the at least one performance parameter for the
30 transaction and updating (328-330) the request by inserting the correlation identifier if the request meets (307-324) the at least one predefined condition being performed under the control of the second module.

3. The method (300) according to claim 1 or 2, further including at least one recursive execution of the steps of:

the server computer originating (344-346) at least one further request of service for at least one further server
5 computer,

if the request includes (348) the correlation identifier, updating (350) the at least one further request by inserting a further correlation identifier,

transmitting (352) each further request to the
10 corresponding further server computer,

if the further request includes (354) the further correlation identifier, enabling (358) the measuring on the further server computer of the at least one performance parameter for a further sub-transaction originating from the
15 further request,

executing (360) the further sub-transaction, and

associating (366) the further correlation identifier with the at least one performance parameter measured on the further server computer.

20 4. The method (300) according to any claim from 1 to 3, wherein the step of enabling (326) the measuring of the at least one performance parameter for the transaction and updating (328-330) the request by inserting the correlation identifier if the request meets (306-324) the at least one
25 predefined condition includes:

verifying (314-316) whether an address of the server computer matches a predetermined pattern stored on the client computer.

5. The method (300) according to any claim from 1 to 4,
30 wherein the step of originating (303-305) the request includes:

downloading (303) a document from the server computer,
displaying (304) the document, and
selecting (305) a link in the document,

and wherein the step of enabling (326) the measuring of the at least one performance parameter for the transaction and updating (328-330) the request by inserting the correlation identifier if the request meets (307-324) the at least one
5 predefined condition includes:

verifying (307-310,318-324) whether a definition of the document includes at least one enabling identifier.

6. The method (300) according to claim 5, wherein the step of verifying (307-310,318-324) whether the definition of the
10 document includes the at least one enabling identifier includes:

verifying (307-310) whether the definition of the document includes a global enabling identifier for all the transactions originating from the document.

15 7. The method (300) according to claim 5 or 6, wherein the step of verifying (307-310,318-324) whether the definition of the document includes the at least one enabling identifier includes:

verifying (322-324) whether a definition of the selected
20 link includes a local enabling identifier for the transactions originating from the selected link.

8. A computer program (205-225;240-260) directly loadable into a working memory (140) of a data processing system (100) with a distributed architecture for performing the method of
25 any claim from 1 to 7 when the program is run on the data processing system.

9. A computer program (205-225), directly loadable into a working memory (140) of a client computer (110) in a data processing system (100) with a distributed architecture, for
30 performing a method (300) of monitoring performance of distributed applications when the program is run on the client computer, the method including the steps of:

originating (303-305) a request of service for a server computer,

if the request meets (307-324) at least one predefined condition, enabling (326) the measuring on the client computer
5 of at least one performance parameter for a transaction corresponding to the request and updating (328-330) the request by inserting a correlation identifier,

transmitting (320) the request to the server computer for causing the server computer to enable (334) the measuring on
10 the server computer of the at least one performance parameter for a sub-transaction originating from the request if the request includes (332) the correlation identifier, to execute (336-372) the sub-transaction, and to associate (378) the correlation identifier with the at least one performance
15 parameter measured on the server computer, and

associating (388) the correlation identifier with the at least one performance parameter measured on the client computer.

10. A computer program (205-225), directly loadable into a
20 working memory (140) of a server computer (120) in a data processing system (100) with a distributed architecture, for performing a method (300) of monitoring performance of distributed applications when the program is run on the server computer, the method including the steps of:

25 receiving (320) a request of service from a client computer, the client computer, if the request meets (307-324) at least one predefined condition, enabling (326) the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and
30 updating (328-330) the request by inserting a correlation identifier, the correlation identifier being associated (388) with the at least one performance parameter measured on the client computer,

if the request includes (332) the correlation identifier,
35 enabling (334) the measuring on the server computer of the at

least one performance parameter for a sub-transaction originating from the request,

executing (336-372) the sub-transaction, and

associating (378) the correlation identifier with the at
5 least one performance parameter measured on the server computer.

11. A computer program (210), directly loadable into a working memory (140) of a client computer (110) in a data processing system (100) with a distributed architecture, for performing a
10 method (300) of monitoring performance of distributed applications when the program is run on the client computer, the method including the steps of:

intercepting (306) a request of service for a server computer originated on the client computer by a further
15 computer program (205),

if the request meets (307-324) at least one predefined condition, enabling (326) the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and updating (328-330) the
20 request by inserting a correlation identifier,

transmitting (320) the request to the server computer for causing the server computer to enable (334) the measuring on the server computer of the at least one performance parameter for a sub-transaction originating from the request if the
25 request includes (332) the correlation identifier, to execute (336-372) the sub-transaction, and to associate (378) the correlation identifier with the at least one performance parameter measured on the server computer, and

associating (388) the correlation identifier with the at
30 least one performance parameter measured on the client computer.

12. A program product (165) comprising a computer readable medium on which the program (205-225;240-260;205) of any claim from 8 to 11 is stored.

13. A data processing system (100) with a distributed architecture for monitoring performance of distributed applications including at least one client computer (110) and at least one server computer (120), wherein each client
5 computer has means (205) for originating a request of service for a server computer, means (210) for enabling the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and for updating the request by inserting a correlation identifier if
10 the request meets at least one predefined condition, means (230) for transmitting the request to the server computer, and means (215,220) for associating the correlation identifier with the at least one performance parameter measured on the client computer, and wherein each server computer has means
15 (255) for enabling the measuring on the server computer of the at least one performance parameter for a sub-transaction originating from the request if the request includes the correlation identifier, means (240,250) for executing the sub-transaction, and means (255,260) for associating the
20 correlation identifier with the at least one performance parameter measured on the server computer.

14. A client computer (110) for monitoring performance of distributed applications in a data processing system (100) with a distributed architecture, the client computer including
25 means (205) for originating a request of service for a server computer, means for (210) enabling the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and for updating the request by inserting a correlation identifier if the request
30 meets at least one predefined condition, means (230) for transmitting the request to the server computer for causing the server computer to enable the measuring on the server computer of the at least one performance parameter for a sub-transaction originating from the request if the request

includes the correlation identifier, to execute the sub-transaction, and to associate the correlation identifier with the at least one performance parameter measured on the server computer, and means (215,220) for associating the
5 correlation identifier with the at least one performance parameter measured on the client computer.

15. A server computer (120) for monitoring performance of distributed applications in a data processing system (100) with a distributed architecture, the server computer including
10 means (235) for receiving a request of service from a client computer, the client computer, if the request meets at least one predefined condition, enabling the measuring on the client computer of at least one performance parameter for a transaction corresponding to the request and updating the
15 request by inserting a correlation identifier, the correlation identifier being associated with the at least one performance parameter measured on the client computer, means (255) for enabling the measuring on the server computer of the at least one performance parameter for a sub-transaction originating
20 from the request if the request includes the correlation identifier, means (240,250) for executing the sub-transaction, and means (255,260) for associating the correlation identifier with the at least one performance parameter measured on the server computer.

**A METHOD AND SYSTEM FOR MONITORING PERFORMANCE OF DISTRIBUTED
APPLICATIONS**

Abstract

5 A method (300) and a corresponding system for monitoring performance of distributed applications are proposed. In the method of the invention, a sensor intercepts (306) every request of service for a server that is generated (303-305) on a client. If the request meets a filtering condition (for
10 example, defined by the address of the server, the web page from which the request is originated and/or the selected link) the measuring of a corresponding transaction on the client is enabled (326); at the same time, the request is updated (328-330) by inserting a correlator. The request is then
15 transmitted (320) to the server. If the request includes the correlator, the measuring of a sub-transaction originating from the request is also enabled (334) on the server. The parameters measured on the client and on the server are then associated (388,378) with the correlator.

20 FIGURES 3a-3d




FIG. 1a




FIG. 1b




FIG. 2




FIG. 3a




FIG. 3b

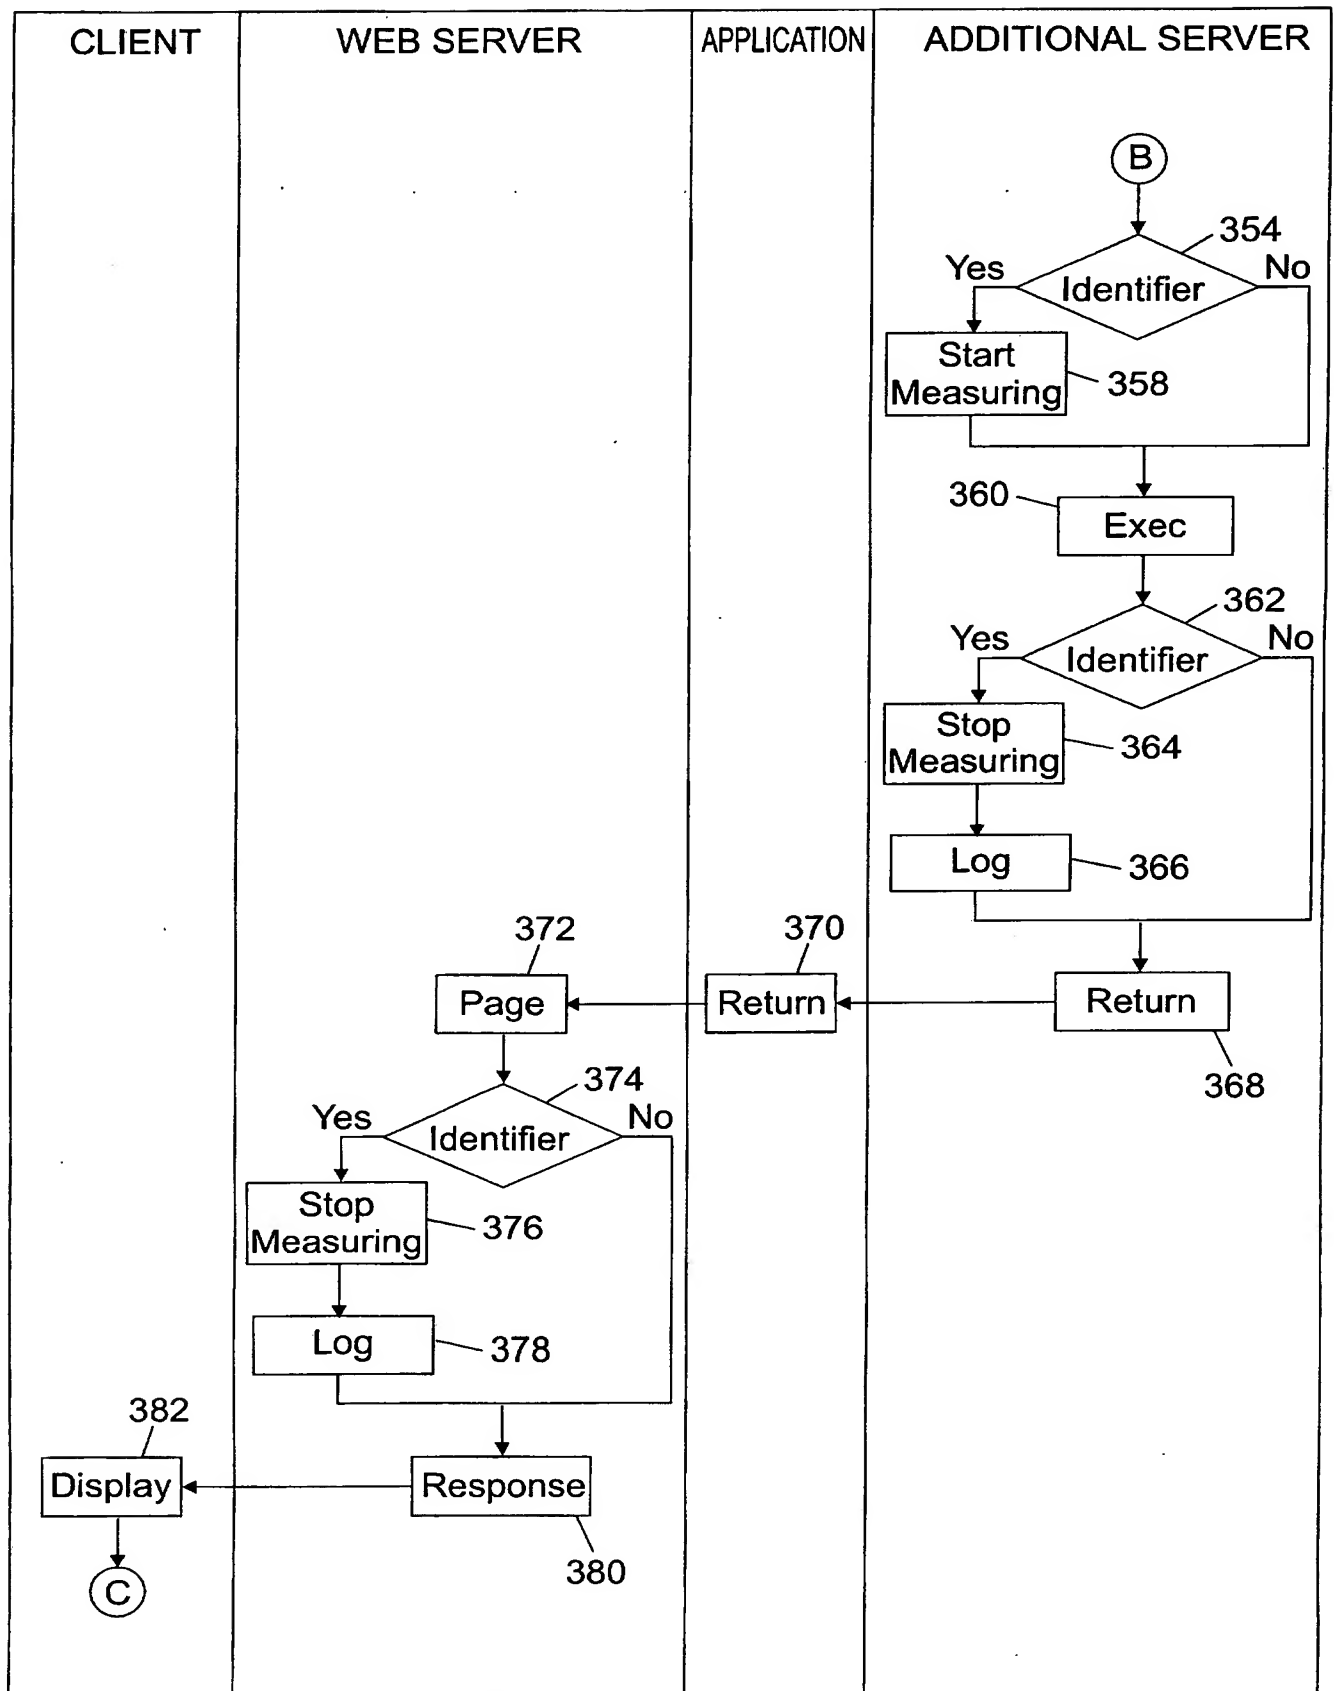


FIG. 3c




FIG. 3d

